

DataCenter 2020: hot aisle and cold aisle containment efficiencies reveal no significant differences

November 2011

Powered by

• • T • • Systems •



DataCenter 2020: hot aisle and cold aisle containment efficiencies reveal no significant differences

The DataCenter 2020 is a joint T-Systems and Intel data center test laboratory in the Munichbased Euroindustriepark. The two companies are working to determine methods by which energy efficiency and operations in existing data centers can be optimized. Based on this work and the knowledge derived, the team is also developing a model of the data center of the future. So far, the researchers have demonstrated improved energy efficiency through the consistent separation of cold and hot air and by raising the supply air temperature and provisioning to a higher energy density in the rack. A substantial proportion of the lower energy consumption is due to being able to optimize the air flow rate through variable speed fans in the cooling units, a capability derived from the cold-aisle containment. In this, the third phase of research, the team has repeated the efforts outlined in the first two white papers, with a hot-aisle containment to compare and contrast it with the improvements found with the cold aisle containment.

Review: The results to date

In the first phase of the optimization the researchers from T-Systems and Intel reduced data center energy consumption by two simple means:

1. The strict separation of cold and hot air in the raised floor and the cold aisle containment lead to optimized air ducting and minimization of leakage and airflow mixing. This can also reduce the fan speed of the circulating air cooling units.
2. Raising the supply air temperature delivered under the raised floor (T1) along with a coincident increase of the chilled water temperature. This minimizes the hours required for cooling by the chillers and extends the hours available for indirect free cooling. The PUE result could be further improved if the cooling air temperature was increased in accordance with the upper limit of the ASHRAE recommendations to provide a computer inlet temperature (TA) of 27 °C.

With these initial measures, the researchers succeeded in reducing the DataCenter 2020 PUE ratio from 1.8 to 1.4, with the computer inlet temperature remaining unchanged at 22 °C. The „Power Usage Effectiveness“ (PUE) value measures the efficiency of energy use in the Data Center infrastructure. It shows how much energy used is actually consumed by the IT equipment versus the data center support equipment.

PUE is the ratio of the total energy used in the data center (Total Facility Power Consumption) and total energy consumption for IT Equipment (IT Equipment Power Consumption). The PUE is a good metric for the energy consumption of the non-IT-related infrastructure (power distribution, airflow, cooling, facility security etc.) supporting the IT equipment. PUE, however, does not offer any conclusions about the energy efficiency of IT equipment or the overall efficiency of the data center itself.




The following figure depicts the changes resulting from reducing the leakage in the raised floor and the cold aisle containment, and the computer inlet temperature (TA) remains constant at 22 °C.

By sealing the raised floor and providing a clear separation of cold and warm air we increased the pressure in the raised floor and increased the ΔT in the CRAH cooling coil (ΔT = difference between return air and supply air temperature) from 6 to 17 °C. This means that the CRAH will operate much more efficiently than before. Additionally, the fan speed was reduced to 30 percent (previously at 100 percent). Because of the fan power laws, this means that the fan motor consumes about 90 percent less energy. The air flow also slows down due to the reduced fan speed. Since the air then absorbs more heat, the return air temperature (T2) is increased from 24 °C to 38 °C. Also, the inlet temperature (T1) under the raised floor can be increased from 18 °C to 21 °C. Since the computer inlet temperature TA is still at 22 °C, a very small temperature gradient between the raised floor and server inlet temperature is achieved. In the next step (see White Paper: Delivering high density in the Data Center; efficiently and reliably), the DataCenter 2020 team increased the IT load to 22kW/rack to see how the PUE or the efficiency of data center infrastructure may depend on the energy density. The computer's inlet temperature (TA) remained constant still at 22 °C.

They chose two scenarios:


- In the first scenario, they used a single CRAH. The external circulation water supply temperature was kept at 8 °C. The resultant PUE decreased at an energy density of 22 kW/rack down to 1.32.
- In the second scenario, they established an energy density of 10kW/rack with two CRAHs operating with a water supply temperature of 16 °C. The two CRAH fan speeds were reduced accordingly. With only one half of the airflow needed with two CRAHs operational, only one quarter of the power was needed in each unit as compared with single unit operation. The higher water-supply temperature also reduces the use of chillers. This also allowed a greater use of indirect free cooling for more hours per year, which also reduces the total energy expenditure. This allowed the team to further reduce the PUE value to 1.23.



Overall, it was found that energy densities well in excess of 20 kW/rack could be supported safely and with good availability using standard technology in conventional CRAH based cooling. The structured use of the optimization steps outlined in our first white paper remains the best path. The cold aisle containment becomes more important at higher energy densities and plays a more important supporting role in failure scenarios. The measurements showed that, with separation of cold and warm air (leakage sealing, cold aisle containment), the energy density or IT load per rack can increase by approximately three times - and yet with the same reliability as in the case without containment. In addition, we benefited from a nearly three-fold higher run time in the event of complete failure of the air supply.

New measures: hot aisle containment



In the next step, the researchers conducted the same measurements, but this time using a hot-aisle containment to examine the preferred performance in comparison to standard configurations (hot aisle/cold aisle) and to compare it with the cold-aisle containment. Different opinions existed in the market on each configuration's efficiency advantages, but nothing had been established empirically. Regardless of the selected enclosure type, if the servers are operated the same, with the same inlet temperature and the airflow path is controlled, the energy consumption and the amount of dissipated heat to be cooled are identical in both systems. For the measurements made with cold-aisle containment, the CRAH operated most efficiently with a temperature difference ΔT of 17 °C (air temperature 38 °C, air temperature 21 °C).



In the hot-aisle containment, the warm exhaust air from the servers is enclosed and routed directly to the CRAH; the space of the data center is the cold aisle. Since the CRAH must provide cold air into a wider area with higher air volume leads the following hypothesis has been offered by others: the hot aisle containment requires a higher fan speed than the cold aisle containment. On the other hand, it has also been proposed that the hot-aisle containment would offer advantages in the CRAH's ΔT (higher efficiency), since the heat or the warm air generated by

the servers is directed through a duct directly to the cooling coil of the CRAH. This warm air is better insulated, air turbulence and mixing reduced, and heat moves directly to the cooler with no losses to any ambient cooling.

This raises the question: are these effects real? Are these the same effects, but just in different sections of each containment scheme? Which system is more efficient?



Basically, three parameters are crucial for the cooling function:

1. First there is the cooling coil of the CRAH, which brings in the warm air and is cooled by cold water as it flows through. The higher the ΔT (= difference between return air and chilled water supply) the better.
2. Secondly, the fan of the CRAH, which defines the flow rate in the data center. Air velocities that are too high can create poor airflow management and even preclude proper distribution into the rack inlets. Further excessive pressure in the cold aisle may result in over-pressurization and blow through the servers, overspinning, with possible detrimental effects on the server fans.
3. Thirdly, airflow management in the room, optimizing airflow under the raised floor, control of flow into the containment and preventing air leakage.

If all of the above-mentioned three parameters are optimized for each case, there should be no differences in the cold aisle and hot aisle containment, since the amount of heat is the same as is the energy used to remove it. This was the assumption of the technical leads in the DataCenter 2020. The remainder of this white paper discusses whether the measurements confirm this hypothesis.

Differences in availability and fan speed, but no differences in PUE

Firstly, the team explored the optimal ratio of energy density and reliability. To do this they recorded temperatures across the data center during a cooling system outage. This was done with different IT densities as well as different spatial configurations. The baseline operation in the DataCenter 2020 was with an inlet temperature to the IT of 22C. The servers deployed (and typical of most IT) had an upper allowable inlet temperature of 35C. The manufacturer typically guarantees operation of the equipment up to that point. The cooling systems were failed with a simulated power outage. Generally the servers will run based on UPS backup. The team's experience has been that data centers typically have a 12-minute UPS capability for the IT equipment. By that time it is typical that backup generators will have come online and will be supplying power to the IT and the cooling system, which will have restarted within this window.

The emergency generators and chillers require a certain amount of time to restart before cooling is returned to the data center to allow continued operation of the IT equipment. The team's goal was to investigate in the three configurations (hot aisle containment, cold aisle containment, and hot/cold aisle arrangements) at a range of power densities per rack, how long before the inlet temperature of the IT equipment reached the upper limit and risked failure of the servers due to overheating ($> 35^{\circ}\text{C}$ inlet temperature).

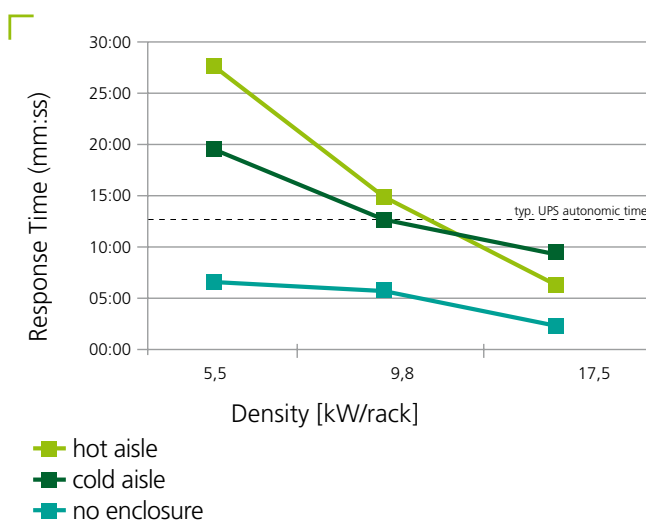

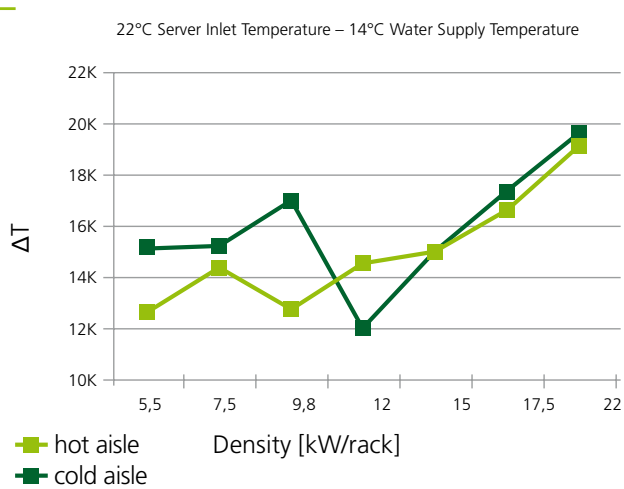



Figure 5 shows that there are definite advantages for hot and cold aisle containment. The low-density racks, particularly those in the hot aisle containment, have far more time before failure from high temperatures. The reason is the large volume of cool air contained in the room. Note that both containment systems have marked advantages as compared with the no-containment system. It has been often suggested, and now shown to be suggested incorrectly, that an open system would allow a longer time before failure. Instead, the lack of containment allows server air recirculation very quickly. At high energy density in the rack of 17.5 kW/rack, the differences appear to be rather small with slight advantages for the cold aisle containment, most likely due to the higher volume of total airflow movement and the cold aisle ability to use the larger high ceiling volume as a heat sink. While the two systems behave differently depending on the energy density in the rack, both cold aisle and hot aisle containments offer greater advantages compared with the scenario without containment. From Figure 5 we see that in the particular case of a rack density of 10kW/rack, both containment systems provide more cooling time than the common IT UPS (>12 minutes), if the cooling capability has not been restored in this time, the IT equipment will be shutting down due to lack of further UPS capability and cooling not required anyway. In no case (rack density) did the open data center provide the time needed to support our target of 12 minutes. Contrary to initial fears, containment systems actually increase server availability as compared with other room layouts.



The same applies to the fan speed in the CRAH, which is already optimized for energy efficiency and running at less than 50 percent. Thus, the bulk of the energy efficiency savings from the CRAH (> 85% energy savings) has already been achieved. Figure 6 also shows other different behaviors. At a lower energy density in the rack the cold aisle containment allowed a lower fan speed, while by increasing the energy density, the lower fan speed was achieved with the hot-aisle containment. The notable jump is explained by the required use of a second circulation cooling device from an energy density of 10kW/rack and above which the airflow dynamics shifted. The team believes that the differences are specific to the system and the site and the resistance of flow in the particular ducting and underfloor in the DC2020, and that these results could well be reversed in a different facility. But in any case, the differences in the optimized fan speeds are extremely low and likely below the detection limit in the electricity bill.



Even when we are looking at ΔT in the CRAH units – within narrow limits – the differences between the two containment systems exist. For lower and higher energy densities, the cold aisle containment seems to have a slight benefit. However, due to the small differences, our conclusion above is that the differences are slight and most likely site-specific related to the exact airflow paths, so there is no consequential difference in efficiency.



In summary, Figure 8 confirms the hypothesis of the researchers at the DataCenter 2020 that there are no differences when operating with optimized parameters for the cold aisle and hot aisle containment, since the amount of heat is always identical. Because there are only minor differences in the value of „Power Usage Effectiveness“ (PUE), which measures the efficiency of the infrastructure’s energy, we take these to be within the measurement tolerance. It is important that the same optimization methodology be consistently applied in both systems; that is, eliminate leakage, separate hot and cold air and eventually reduce the fan speed of the recirculation cooling unit. Differences may exist in efficiency if these are not consistently applied. However, the energy efficiency has not then been optimized and we would not expect to be able to compare the results.

Criteria for selection of the system

Once it is seen that there are no differences in cold and aisle containment systems relating to energy efficiency, the choice of which to use becomes one of an operational or space architectural view. A look at the advantages and disadvantages of the cold aisle and hot aisle containment systems will make the system choice easier.

The cold aisle containment system can separate cold air and warm air from each other, without major modification of the server room. This may be as simple as a false ceiling over the racks and doors at the ends of rows. It can also be done with a low ceiling height. Gas-based fire extinguishing systems can also be used effectively as the quenching gas can flow freely through the raised floor and into the containment. However, the down side of the cold aisle containment is that employees must work in the open warm space which could be about 45 °C or warmer, even though the cold air supplied to the IT equipment continues to be optimal. In addition, a raised floor for air circulation is typically needed to transport cool air into the cold aisle.

As with the cold aisle containment, the hot aisle containment also separates the cold air and warm air from each other. The main advantage for the hot-aisle containment is for personnel. The hot aisle containment creates a pleasant working atmosphere for personnel as they walk into a cool room, (even if much of the cabling work is still in the hot aisle). In addition, a hot-aisle containment structure can be integrated into an existing space in order to eliminate „hot spots“ without putting a burden on the remaining room environment with excess heat re-distribution. In principle, from an airflow perspective, there is no need for a raised floor. They are well suited to new construction. On the down side there are the possible investments required in the reconstruction of the gas extinguishing system, and additional structures needed for the air ducting. If there is no raised floor, all supply lines (power and network cabling) must be from above. This can lead to interference with the airflow path. Therefore, the design of a data center with hot-aisle containment may still consider a raised floor (assuming sufficient building height). In both cases, and with lower airflow mixing, there will be areas much warmer than in a typical open data center. Therefore all materials and components must be verified as suitable for the hottest expected temperatures that they may be exposed to. The same must be done for the safety and well-being of the employees. Operating temperatures where people will be working must be understood and designed for.

Both systems have their advantages and disadvantages, but they are equally effective at providing efficient well-managed airflow to the IT equipment. Therefore, the choice depends on the particular requirements and operating conditions and architectural features of the individual site.

Conclusion

The researchers at the DataCenter 2020 improved the center's PUE; with standard equipment and a range of energy efficiency optimization measures. We achieved a PUE of 1.23 without changing the environmental conditions for the IT equipment. This was achieved by a cold or hot-aisle containment program (there are no efficiency differences between the two), a lower CRAH fan speed, and a higher energy density in the server rack. The enclosure of the cold or hot aisles increases the availability and reliability of the servers. Thus, an enclosure at energy densities of 10 to 17.5 kW/rack can be an interesting alternative and relatively inexpensive compared with a larger UPS for air conditioning, although still not being able to replace other availability measures completely.

The best result for PUE was achieved when the team operated the data center according to the upper limit of the ASHRAE recommendations, with a computer inlet temperature (TA) of 27 °C. The air temperature below the raised floor (T1) is then at 26 °C. This then allows the water supply from the chiller to rise to 20 °C. Experience with the CRAH cooling coil has shown that we needed a water temperature that is usually around 6 °C lower than the temperature of the air under the floor. The higher water temperature reduces the time required for cooling using the chillers and extends the amount of operating time for indirect free cooling by outside air.

Outlook: chiller-less cooling in favorable climates reduces investment costs

The amount of hours of indirect free cooling with outside air is dependent on climatic conditions at the location of a data center. Take, for example Munich, the site of the DataCenter 2020:

providing IT equipment with an inlet temperature of 27 °C requires a raised-floor air temperature of 26 °C, which in turn means that the cold water supply has to be at around 20 °C.

This water temperature can be achieved comfortably by indirect free cooling as long as the outside air temperature does not exceed 20 °C (around 10% of the hours each year in Munich, see Figure 9). Furthermore, a hybrid dry/wet cooling solution can achieve such water temperature levels with an outside air temperature of up to 26 °C by utilizing the effect of evaporation of treated water, a very efficient overall method.




Figure 9

With a combination of measures like this, the energy-intensive method of compression cooling will be required on only four days during an average year in Munich, allowing a general rethinking of cooling system philosophies for future data centers.

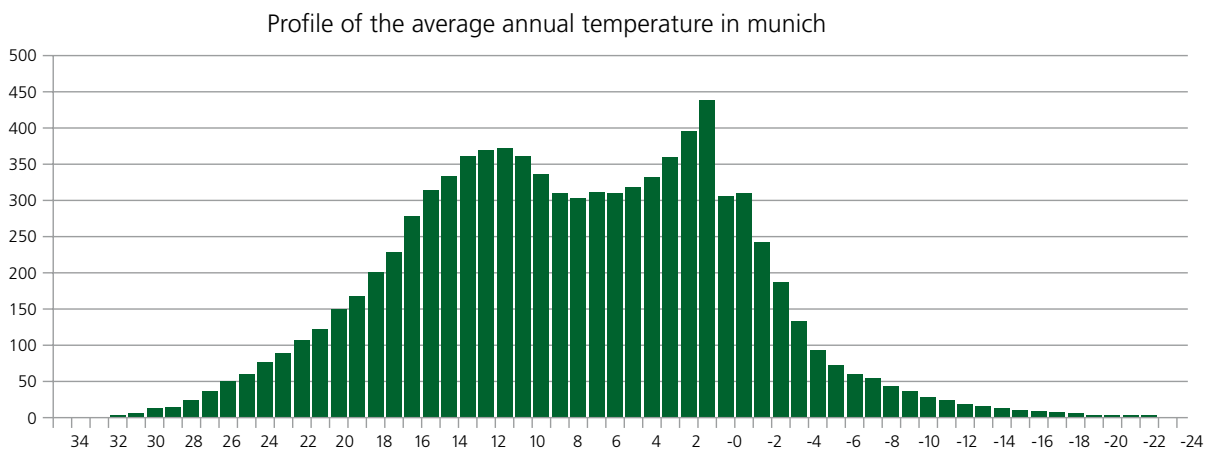


Figure 10

For more information on the DataCenter 2020, see: www.datacenter2020.com

Copyright © 2010 Intel Corporation and T-Systems International GmbH
All rights reserved.

Powered by

